AMENDMENTS TO THE CLAIMS

This listing of claims will replace all prior versions, and listings, of claims in the application:

Listing of Claims:

1. (Previously Presented) A method of flow controlling InfiniBand receive traffic, comprising:

maintaining a single memory structure for queuing InfiniBand traffic received via multiple virtual lanes and multiple queue pairs;

identifying a first packet payload received via a first virtual lane and a first queue pair;

determining whether the first payload can be stored in the memory structure without exceeding a portion of the memory structure allocated to the first virtual lane;

determining whether the first payload can be stored in the memory structure without exceeding a portion of the memory structure allocated to the first queue pair;

if storing the first payload in the memory structure would exceed said portion of the memory structure allocated to the first queue pair, determining whether the first queue pair is enabled to use a shared portion of the memory structure to store payloads of packets received via the first queue pair; and

maintaining a second memory configured to store, for each of the multiple queue pairs that is active, one or more parameters associated with operation of said queue pair, wherein said parameters include:

a maximum number of message credits advertisable by said queue pair;

a maximum number of memory structure buffers dedicated to storing payloads of packets received via said queue pair;

an indicator configured to indicate whether said queue pair is enabled to use a set of shared memory structure buffers; and

a number of shared memory structure buffers in said set of shared memory structure buffers, wherein said shared memory structure buffers are available for use by said queue pair to store payloads of packets received via said queue pair if:

said queue pair has used said maximum number of memory structure buffers; and

said indicator indicates that said queue pair is enabled to use said set of shared memory structure buffers; and

a maximum number of message credits advertisable by said queue pair when said queue pair starts using said shared memory structure buffers.

2-10. (Cancelled)

11. (Previously Presented) A computer readable medium storing instructions that, when executed by a computer, cause the computer to perform a method of flow controlling InfiniBand receive traffic, the method comprising:

maintaining a single memory structure for queuing InfiniBand traffic received via multiple virtual lanes and multiple queue pairs;

identifying a first packet payload received via a first virtual lane and a first queue pair;

determining whether the first payload can be stored in the memory structure without exceeding a portion of the memory structure allocated to the first virtual lane;

determining whether the first payload can be stored in the memory structure without exceeding a portion of the memory structure allocated to the first queue pair;

if storing the first payload in the memory structure would exceed said portion of the memory structure allocated to the first queue pair, determining whether the first queue pair is enabled to use a shared portion of the memory structure to store payloads of packets received via the first queue pair; and

maintaining a second memory configured to store, for each of the multiple queue pairs that is active, one or more parameters associated with operation of said queue pair, wherein said parameters include:

a maximum number of message credits advertisable by said queue pair;

a maximum number of memory structure buffers dedicated to storing payloads of packets received via said queue pair;

an indicator configured to indicate whether said queue pair is enabled to use a set of shared memory structure buffers; and

a number of shared memory structure buffers in said set of shared memory structure buffers, wherein said shared memory structure buffers are available for use by said queue pair to store payloads of packets received via said queue pair if:

said queue pair has used said maximum number of memory structure buffers; and

said indicator indicates that said queue pair is enabled to use said set of shared memory structure buffers; and

a maximum number of message credits advertisable by said queue pair when said queue pair starts using said shared memory structure buffers.

12-29. (Cancelled)

30. (Previously Presented) A method of avoiding locking, in receive InfiniBand queues, the method comprising:

maintaining a single memory structure for reassembling InfiniBand traffic received via multiple virtual lanes and multiple queue pairs;

identifying a first packet payload received via a first queue pair that is idle, wherein the first queue pair is considered idle if no traffic from the first queue pair is stored in said single memory structure;

for each other queue pair for which traffic from said queue pair is stored in said single memory structure, determining whether sufficient space in the single memory structure is reserved for reassembling said traffic;

storing the first packet payload in said single memory structure only if sufficient space in the single memory structure is available for reassembling said traffic; and

maintaining a second memory configured to store, for each of the multiple queue pairs that is active, one or more parameters associated with operation of said queue pair, wherein said parameters include:

a maximum number of message credits advertisable by said queue pair;

a maximum number of memory structure buffers dedicated to storing payloads of packets received via said queue pair;

an indicator configured to indicate whether said queue pair is enabled to use a set of shared memory structure buffers; and

a number of shared memory structure buffers in said set of shared memory structure buffers, wherein said shared memory structure buffers are available for use by said queue pair to store payloads of packets received via said queue pair if:

said queue pair has used said maximum number of memory structure buffers; and

said indicator indicates that said queue pair is enabled to use said set of shared memory structure buffers; and

a maximum number of message credits advertisable by said queue pair when said queue pair starts using said shared memory structure buffers.

31. (Cancelled)

32. (Previously Presented) An apparatus for flow controlling received InfiniBand traffic, comprising:

a single memory structure configured to queue payloads of InfiniBand traffic received via multiple virtual lanes and multiple queue pairs;

a resource manager configured to manage the memory structure;

a first module configured to facilitate the advertisement of virtual lane credits;

a second module configured to facilitate the advertisement of queue pair credits; and

a second memory configured to store, for each of the multiple queue pairs that is active, one or more parameters associated with operation of said queue pair, wherein said parameters include:

a maximum number of message credits advertisable by said queue pair;

a maximum number of memory structure buffers dedicated to storing payloads of packets received via said queue pair;

an indicator configured to indicate whether said queue pair is enabled to use a set of shared memory structure buffers; and

a number of shared memory structure buffers in said set of shared memory structure buffers, wherein said shared memory structure buffers are available for use by said queue pair to store payloads of packets received via said queue pair if:

said queue pair has used said maximum number of memory structure buffers; and

said indicator indicates that said queue pair is enabled to use said set of shared memory structure buffers; and

a maximum number of message credits advertisable by said queue pair when said queue pair starts using said shared memory structure buffers.

33-55. (Cancelled)

56. (New) A method of flow controlling InfiniBand receive traffic, comprising:

maintaining a single memory structure for queuing InfiniBand traffic received via multiple virtual lanes and multiple queue pairs;

maintaining a second memory configured to store, for each of the multiple queue pairs that is active, one or more parameters associated with operation of said queue pair, wherein said parameters include a maximum number of message credits advertisable by said queue pair;

identifying a first packet payload received via a first virtual lane and a first queue pair;

determining whether the first payload can be stored in the memory structure without exceeding a portion of the memory structure allocated to the first virtual lane;

determining whether the first payload can be stored in the memory

structure without exceeding a portion of the memory structure allocated to the first queue pair; and

if storing the first payload in the memory structure would exceed said portion of the memory structure allocated to the first queue pair, determining whether the first queue pair is enabled to use a shared portion of the memory structure to store payloads of packets received via the first queue pair.

57. (New) The method of claim 56, wherein the second memory is further configured to store one or more additional parameters associated with operation of said queue pair, wherein said additional parameters include:

a maximum number of memory structure buffers dedicated to storing payloads of packets received via said queue pair;

an indicator configured to indicate whether said queue pair is enabled to use a set of shared memory structure buffers; and

a number of shared memory structure buffers in said set of shared memory structure buffers, wherein said shared memory structure buffers are available for use by said queue pair to store payloads of packets received via said queue pair if:

said queue pair has used said maximum number of memory structure buffers; and

said indicator indicates that said queue pair is enabled to use said set of shared memory structure buffers; and

a maximum number of message credits advertisable by said queue pair when said queue pair starts using said shared memory structure buffers.

58. (New) The method of claim 56, further comprising:

allocating a portion of the memory structure to each of the multiple virtual lanes; and

allocating a portion of the memory structure to each of the multiple queue pairs.

59. (New) The method of claim 56, wherein the memory structure comprises a set of linked lists of memory structure buffers, including one linked list for each of the multiple queue pairs that are active.

60. (New) The method of claim 56, further comprising:

dropping the first payload if the first payload cannot be stored in the memory structure without exceeding the portion of the memory structure allocated to the first virtual lane.

61. (New) The method of claim 56, further comprising:

issuing a Retry, Not Ready, Negative Acknowledgement (RNR-NAK) if:

the first payload cannot be stored in the memory structure without exceeding a portion of the memory structure allocated to the first queue pair; and

the first queue pair is not enabled to use the shared portion of the memory structure.

62. (New) The method of claim 56, further comprising:

issuing a Retry, Not Ready, Negative Acknowledgement (RNR-NAK) if:

the first payload cannot be stored in the memory structure without exceeding a portion of the memory structure allocated to the first queue pair;

the first queue pair is enabled to use the shared portion of the memory structure; and

the shared portion of the memory structure is full.

63. (New) The method of claim 56, further comprising:

defining one or more dedicated thresholds in the portion of the memory structure allocated to the first queue pair; and

for each of said dedicated thresholds, identifying a number of message credits the queue pair may advertise when the amount of the memory structure used by the queue pair exceeds said dedicated threshold.

64. (New) The method of claim 56, further comprising:

defining one or more shared thresholds in the shared portion of the memory structure; and

for each of said shared thresholds, identifying a number of message credits the queue pair may advertise when the amount of the shared portion used by the multiple queue pairs exceeds said shared threshold.

65. (New) The method of claim 56, further comprising:

receiving a request on a second queue pair to perform an RDMA (Remote Direct Memory Access) Read operation; and

based on an amount of data expected to be received via the RDMA Read operation, reserving a sufficient number of buffers in the memory structure.

66. (New) The method of claim 56, further comprising:

in the single memory structure, reassembling the queued InfiniBand traffic into outbound communications;

receiving a payload on an idle queue pair, wherein a queue pair is idle if no traffic from the queue pair is stored in the single memory structure; and

only queuing the payload in the single memory structure if sufficient space in the single memory structure is reserved for completing reassembly of outbound communications on each non-idle queue pair.

67. (New) A computer readable medium storing instructions that, when executed by a computer, cause the computer to perform a method of flow controlling InfiniBand receive traffic, the method comprising:

maintaining a single memory structure for queuing InfiniBand traffic received via multiple virtual lanes and multiple queue pairs;

maintaining a second memory configured to store, for each of the multiple queue pairs that is active, one or more parameters associated with operation of said queue pair, wherein said parameters include a maximum number of message credits advertisable by said queue pair;

identifying a first packet payload received via a first virtual lane and a first queue pair;

determining whether the first payload can be stored in the memory structure without exceeding a portion of the memory structure allocated to the first virtual lane;

determining whether the first payload can be stored in the memory structure without exceeding a portion of the memory structure allocated to the first queue pair; and

if storing the first payload in the memory structure would exceed said portion of the memory structure allocated to the first queue pair, determining whether the first queue pair is enabled to use a shared portion of the memory structure to store payloads of packets received via the first queue pair.

68. (New) The computer readable medium of claim 67, wherein the second memory is further configured to store one or more additional parameters associated with operation of said queue pair, wherein said additional parameters include:

a maximum number of memory structure buffers dedicated to storing payloads of packets received via said queue pair;

an indicator configured to indicate whether said queue pair is enabled to use a set of shared memory structure buffers; and

a number of shared memory structure buffers in said set of shared memory structure buffers, wherein said shared memory structure buffers are available for use by said queue pair to store payloads of packets received via said queue pair if:

said queue pair has used said maximum number of memory structure buffers; and

said indicator indicates that said queue pair is enabled to use said set of shared memory structure buffers; and

a maximum number of message credits advertisable by said queue pair when said queue pair starts using said shared memory structure buffers.

69. (New) The computer readable medium of claim 67, wherein the method further comprises:

defining one or more dedicated thresholds in the portion of the memory structure allocated to the first queue pair; and

for each of said dedicated thresholds, identifying a number of message credits the queue pair may advertise when the amount of the memory structure used by the queue pair exceeds said dedicated threshold.

70. (New) The computer readable medium of claim 67, wherein the method further comprises:

defining one or more shared thresholds in the shared portion of the memory structure; and

for each of said shared thresholds, identifying a number of message credits the queue pair may advertise when the amount of the shared portion used by the multiple queue pairs exceeds said shared threshold.

71. (New) The computer readable medium of claim 67, wherein the method further comprises issuing a Retry, Not Ready, Negative Acknowledgement (RNR-NAK) only if one of:

(a) the first payload cannot be stored in the memory structure without exceeding a portion of the memory structure allocated to the first queue pair; and

the first queue pair is not enabled to use the shared portion of the memory structure; and

(b) the first payload cannot be stored in the memory structure without exceeding a portion of the memory structure allocated to the first queue pair;

the first queue pair is enabled to use the shared portion of the memory structure; and

the shared portion of the memory structure is full.

72. (New) A method of avoiding locking in receive InfiniBand queues, the method comprising:

maintaining a single memory structure for reassembling InfiniBand traffic received via multiple virtual lanes and multiple queue pairs;

maintaining a second memory configured to store, for each of the multiple queue pairs that is active, one or more parameters associated with operation of said queue pair, wherein said parameters include a maximum number of message credits advertisable by said queue pair;

identifying a first packet payload received via a first queue pair that is idle, wherein the first queue pair is considered idle if no traffic from the first queue pair is stored in said single memory structure;

for each other queue pair for which traffic from said queue pair is stored in said single memory structure, determining whether sufficient space in the single

memory structure is reserved for reassembling said traffic; and

storing the first packet payload in said single memory structure only if sufficient space in the single memory structure is available for reassembling said traffic.

73. (New) The method of claim 72, wherein the second memory is further configured to store one or more additional parameters associated with operation of said queue pair, wherein said additional parameters include:

a maximum number of memory structure buffers dedicated to storing payloads of packets received via said queue pair;

an indicator configured to indicate whether said queue pair is enabled to use a set of shared memory structure buffers; and

a number of shared memory structure buffers in said set of shared memory structure buffers, wherein said shared memory structure buffers are available for use by said queue pair to store payloads of packets received via said queue pair if:

said queue pair has used said maximum number of memory structure buffers; and

said indicator indicates that said queue pair is enabled to use said set of shared memory structure buffers; and

a maximum number of message credits advertisable by said queue pair when said queue pair starts using said shared memory structure buffers.

74. (Original) The method of claim 72, wherein said determining comprises, for each said other queue pair:

identifying an amount of space in said single memory structure reserved for said other queue pair; and

comparing said amount of reserved space to an amount of space expected

to be needed to complete reassembly of said traffic from said other queue pair.

75. (New) An apparatus for flow controlling received InfiniBand traffic, comprising:

a single memory structure configured to queue payloads of InfiniBand traffic received via multiple virtual lanes and multiple queue pairs;

a second memory configured to store, for each of the multiple queue pairs that is active, one or more parameters associated with operation of said queue pair, wherein said parameters include a maximum number of message credits advertisable by said queue pair;

a resource manager configured to manage the memory structure;

a first module configured to facilitate the advertisement of virtual lane credits;

a second module configured to facilitate the advertisement of queue pair credits.

76. (New) The apparatus of claim 75, wherein the second memory is further configured to store one or more additional parameters associated with operation of said queue pair, wherein said additional parameters include:

a maximum number of memory structure buffers dedicated to storing payloads of packets received via said queue pair;

an indicator configured to indicate whether said queue pair is enabled to use a set of shared memory structure buffers; and

a number of shared memory structure buffers in said set of shared memory structure buffers, wherein said shared memory structure buffers are available for use by said queue pair to store payloads of packets received via said queue pair if:

said queue pair has used said maximum number of memory structure buffers; and

said indicator indicates that said queue pair is enabled to use said set of shared memory structure buffers; and

a maximum number of message credits advertisable by said queue pair when said queue pair starts using said shared memory structure buffers.

77. (New) The apparatus of claim 75, wherein said single memory structure comprises multiple linked lists of memory structure buffers, including one linked list for each of the multiple queue pairs that is active.

78. (New) The apparatus of claim 75, wherein said first module comprises an InfiniBand link core.

79. (New) The apparatus of claim 75, wherein said second module comprises an acknowledgement generator configured to generate transport layer acknowledgements.

80. (New) The apparatus of claim 75, further comprising a processor interface configured to facilitate the programming of operating parameters associated with the multiple virtual lanes and the multiple queue pairs.

81. (New) The apparatus of claim 75, further comprising:

a first memory configured to store one or more parameters associated with operation of a first virtual lane.

82. (New) The apparatus of claim 81, wherein said one or more parameters include:

a count of the number of memory structure buffers currently used to store

payloads of packets received via the first virtual lane; and

a threshold, wherein a first packet is dropped if storing the payload of the first packet would cause said count to exceed said threshold.

83. (New) The apparatus of claim 75, wherein said one or more parameters further include:

one or more dedicated thresholds, wherein each said dedicated threshold identifies a subset of said maximum number of memory structure buffers; and

for each said dedicated threshold, a number of message credits advertisable by said queue pair when said queue pair uses said subset of said maximum number of memory structure buffers.

84. (New) The apparatus of claim 75, wherein said one or more parameters further include:

one or more shared thresholds, wherein each said shared threshold identifies a subset of said number of shared memory structure buffers; and

for each said shared threshold, a number of message credits advertisable by said queue pair when said queue pair uses said subset of said number of shared memory structure buffers.